

Package ‘npwbs’

July 28, 2026

Type Package

Title Nonparametric Multiple Change Point Detection Using Wild Binary Segmentation

Version 0.4.0

Author Gordon J. Ross [aut, cre]

Maintainer Gordon J. Ross <gordon.ross@ed.ac.uk>

Description Implements nonparametric multiple change-point detection for univariate sequences using Wild Binary Segmentation, as described in Ross (2026) “Nonparametric Detection of Multiple Location-Scale Change Points via Wild Binary Segmentation” [doi:10.48550/arXiv.2107.01742](https://doi.org/10.48550/arXiv.2107.01742). The package provides Mann--Whitney, Mood, Lepage, Cramér--von Mises, and modified Baumgartner rank-based statistics, together with method-specific thresholds for controlling the probability of incorrectly detecting a change point in a homogeneous sequence.

Depends R (>= 3.6.0)

Imports Rcpp

LinkingTo Rcpp

License GPL-3

Encoding UTF-8

NeedsCompilation yes

Repository CRAN

Date/Publication 2026-07-28 08:20:02 UTC

Contents

detectChanges	2
Index	4

detectChanges	<i>Nonparametric detection of multiple change points using Wild Binary Segmentation</i>
---------------	---

Description

Returns the estimated number and locations of the change points in a sequence of univariate observations. For full details of how this procedure works, please see G. J. Ross (2026) - "Non-parametric Detection of Multiple Location-Scale Change Points via Wild Binary Segmentation" at <https://arxiv.org/abs/2107.01742>

Usage

```
detectChanges(y, alpha=0.05, prune=TRUE, M=10000, d=2,
              displayOutput=FALSE, method="lepage", breakTies=TRUE)
```

Arguments

y	The sequence to test for change points
alpha	Required Type I error rate. Lepage supports 0.05 and 0.01; the other methods support 0.05.
prune	Whether to prune potential excess change points via post-processing. Most likely should be left as TRUE.
M	Number of subsequences to sample during WBS and pruning. The built-in thresholds are calibrated for $M = 10000$, so other values are not supported by the packaged calibration.
d	Endpoint trimming within each sampled interval. The built-in thresholds support only $d = 2$.
displayOutput	If TRUE then will print some information while searching for change points
method	Rank statistic: exactly one of "lepage", "mw", "mood", "cvm", or "baumgartner".
breakTies	If TRUE, tied inputs receive one random strict ordering before detection; if FALSE, ties are an error.

Details

The built-in thresholds are provided for Lepage, Mann–Whitney, Mood, Cramér–von Mises, and modified Baumgartner statistics using $d = 2$ and $M = 10000$. Lepage at $\alpha = 0.05$ uses the final registered book-chapter calibration, while Lepage at $\alpha = 0.01$ retains the established package calibration. Mann–Whitney, Mood, Cramér–von Mises, and modified Baumgartner support $\alpha = 0.05$ only. The sampled interval minimum length is fixed at 10. Thresholds are provided through segment length $n = 10000$; longer inputs reuse the $n = 10000$ threshold with one warning. Modified-Baumgartner values above $n = 3000$ use the fitted extension accepted for the package rather than direct Monte-Carlo calibration.

If `prune = TRUE`, a post-processing pruning step retests merged neighbouring segments and removes changepoints that are no longer supported.

Value

A vector containing the detected changepoint locations. A returned value k denotes a split between observations k and $k + 1$, i.e. the change occurs after observation k .

Author(s)

Gordon J. Ross <gordon.ross@ed.ac.uk>

Examples

```
set.seed(100)
y <- c(rnorm(30,0,1), rnorm(30,3,1), rnorm(30,0,1), rnorm(30,0,3))
detectChanges(y)
```

Index

`detectChanges`, [2](#)